RESPONSE ALLOCATION IN CONCURRENT CHAINS WHEN TERMINAL-LINK DELAYS FOLLOW AN ASCENDING AND DESCENDING SERIES

DARREN R. CHRISTENSEN

UNIVERSITY OF ARKANSAS FOR MEDICAL SCIENCES

AND

RANDOLPH C. GRACE

UNIVERSITY OF CANTERBURY

Eight pigeons were trained in a concurrent-chains procedure in which the terminal-link immediacy ratio followed an ascending or descending series. Across sessions, one terminal-link delay changed from 2 s to 32 s to 2 s or from 32 s to 2 s to 32 s, while the other was always 8 s. For all pigeons, response allocation tracked changes in delay and was biased towards the 8-s alternative on the descending series, indicating a hysteresis effect, and was more sensitive to changes in the terminal-link delay ratio for relatively long (> 8 s) than short (< 8 s) delays. Both the hysteresis and effect of delay duration were predicted by an extended version of Grace and McLean's (2006) decision model. The extended decision model provided an overall better account of the results than a simple linear-operator model (Grace, 2002), and holds promise for an integrated account of choice in concurrent chains for both acquisition and steady-state conditions.

Key words: reinforcer delay, acquisition, hysteresis, terminal-link effect, concurrent chains, choice, pigeons

---

Traditional research on behavioral choice has used steady-state designs in which subjects are trained with a particular set of contingencies until response allocation stabilizes (e.g., Herrnstein, 1961; see Davison & McCarthy, 1988, for review). A variable such as the relative rate or immediacy of reinforcement across the alternatives is then varied parametrically across conditions. Results from these experiments are typically well described by models based on the matching law, which in its most general form states that response allocation matches the relative value obtained from the choice alternatives (Baum & Rachlin, 1969). For concurrent chains, such models include delay-reduction theory (Fantino, Preston & Dunn, 1993), the contextual choice model (Grace, 1994), and the hyperbolic-value added model (Mazur, 2001). These models differ in terms of specific details, but all share the assumption that initial-link response allocation in concurrent chains matches the relative value associated with the terminal links.

However, there is a growing literature on acquisition of choice—how response alloca-

tion changes when the reinforcement contingencies are altered (e.g., Davison & Baum, 2000; Mazur, 1992, 1995, 1996; Mazur, Blake, & McManus, 2001). An important question is whether the principles that describe choice at the steady-state level—such as the assumption that response allocation matches relative value—also apply to choice in transition. For example, Grace (2002) trained pigeons on a concurrent-chains procedure in which the location of the shorter terminal-link schedule was changed every 20 sessions. Across conditions, he studied transitions between different combinations of terminal-link schedules. He found that acquisition of preference was well described by a simple linear-operator model (LINOP). The LINOP model incorporated the basic assumption of the matching law, that is, that response allocation matched the relative value of the terminal-link schedules. Also, the asymptotic value of a schedule (i.e., after sufficient exposure to the schedule) was defined as a function of the distribution of delays to reinforcement (Shull, Spear, & Bryson, 1981)—which is a common assumption for models of steady-state choice (cf. Mazur, 1984, 2001; Grace, 1996). Third, the model assumed that when the terminal links were changed, value was updated according to

a linear-operator rule:

$$\Delta V_{n+1} = r\left(V_{asymp} - V_n\right) \qquad (1)$$

According to Equation 1, the change in value for cycle n+1 is a constant proportion of the difference between the asymptotic value and the value on cycle n. Grace (2002) showed that the LINOP model made more accurate predictions than a competing memory-representational model. The important point to emphasize is that the LINOP model is based on assumptions which are common to steady-state models: matching to relative value, with value determined as a function of the reinforcer delay distribution.

Several studies on choice acquisition have used a procedure originally devised by Hunter and Davison (1985), in which reinforcement contingencies change unpredictably across sessions according to a pseudorandom binary series (PRBS). Effectively, a PRBS series ensures that the reinforcer ratio in the current session cannot be predicted from those in prior sessions. This research has shown that choice responding can adjust very rapidly to changes in reinforcement contingencies, and stimulated the development of models which are not explicitly derived from steady-state accounts of choice.

For example, Schofield and Davison (1997) showed that pigeons' response allocation in concurrent variable-interval (VI) VI schedules tracked changes in the reinforcer ratio when ratios changed daily according to a 31-step PRBS. They conducted a multiple regression analysis and showed that the coefficient that measured sensitivity to the reinforcer ratio was significant and positive for the current session (i.e., Lag 0), but was not significant for the preceding nine sessions (Lag 1 through Lag 9), after three PRBS presentations (93 sessions). Thus, response allocation was controlled by the reinforcer ratio arranged in the current session with virtually no effect from prior sessions. Because cumulative sensitivity levels were similar to those obtained in past research (Baum, 1979), Schofield and Davison suggested that the PRBS design might present an attractive alternative to steady-state designs. However, their procedure is potentially even more important in terms of providing a rich dataset—an acquisition curve

in each session—which presumably reflects the same response-generating process that determines choice in steady-state designs. If so, it is possible that understanding how response allocation adapts to a variable environment may provide insights into steady-state phenomena such as matching.

Grace, Bragason, and McLean (2003) applied the PRBS design to the concurrent-chains procedure to study acquisition of choice between delayed reinforcers. In concurrent chains, subjects respond during a choice phase (initial links) to produce one of two mutually-exclusive outcome schedules which end with reinforcer delivery (terminal links). The relative reinforcer immediacy during the terminal links (i.e., ratio of the reciprocal of the reinforcer delays) is a major determiner of response allocation during the initial links; Grace (1994) showed that an extension of the generalized matching law (Baum, 1974; Davison, 1983) that assumes subjects' relative initial-link response rates match the relative value of the terminal-link schedules, with value determined as a power function of the immediacy ratio, provides an excellent account of response allocation in concurrent chains (cf. Mazur, 2001).

In Grace et al.'s (2003) Experiment 1, the terminal-link schedule associated with the left alternative was always fixed interval (FI) 8 s, while the right terminal-link schedule changed between FI 4 s or FI 16 s across sessions according to a 31-step PRBS. Grace et al. conducted a multiple regression analysis similar to Schofield and Davison's (1997) and found that initial-link response allocation was most sensitive to the immediacy ratio in the current session. The average Lag 0 sensitivity coefficient was 1.04, and varied between 0.47 and 1.84 across subjects. Although these values are lower than those generally obtained in steady-state research (see Grace, 1994, for review), Grace et al.'s results show that response allocation tracked unpredictable daily changes in the terminal-link immediacy ratio.

The same subjects served in Grace et al.'s (2003) Experiment 2, in which a different value for the right terminal link FI schedule was used in each session while the left terminal link remained FI 8 s. Schedule values for the right terminal link varied between 2 s and 32 s, and were determined pseudorandomly such that the log immediacy ratios were uniformly

distributed between log(1/4) and log(4), with the location of the shorter terminal link for each session determined by the PRBS. Thus, the average log immediacy ratio for sessions in which the shorter terminal link was associated with the right (or left) alternative was the same as in Experiment 1 (i.e., log[1/2] or log[2]). Grace et al. found that sensitivity to immediacy did not differ systematically from Experiment 1, suggesting that whether the changing terminal-link schedule took either two (Experiment 1) or a potentially unlimited number of values (Experiment 2) did not affect sensitivity to the immediacy ratio. Also interesting was that for one pigeon the relationship between the log initial-link response ratio and the log immediacy ratio for the current session (as shown in a generalized-matching scatterplot) was nonlinear, with data points falling into two clusters. Grace et al. suggested that a process similar to categorical discrimination might have determined responding for this pigeon.

Grace and McLean (2006) provided a stronger test of whether the degree of variation in delays affected sensitivity to immediacy in concurrent chains when the position of the richer terminal-link changed across sessions according to a PRBS. They compared sensitivity in a ''minimal variation'' condition that was identical to Grace et al.'s (2003) Experiment 1 in which one terminal link was constant (FI 8 s) while the other was either FI 4 s or FI 16 s, with a ''maximal variation'' condition in which a different pair of terminal-link schedules was used in every session. In both conditions, the average log immediacy ratio for sessions in which the richer terminal link was associated with the left (or right) key was log(2) (or log[1/2]). Each condition consisted of three PRBS presentations (93 sessions), and the order was counterbalanced. They found that response allocation tracked the current-session immediacy ratio in both conditions, but that across subjects there was no systematic difference in sensitivity to immediacy. Additionally, in the maximal-variation condition the scatterplot of the log initial-link response ratio as a function of the log immediacy ratio was nonlinear (sigmoidal) for one pigeon in the third PRBS presentation and for a second pigeon when the condition was replicated, again suggesting a categorical discrimination. However, in other cases scatterplots were approximately linear (including the third

PRBS presentation for the subject whose data were nonlinear in the replication), consistent with a traditional generalized-matching model. The sigmoidal relation in generalized-matching scatterplots for pigeons responding under rapid-acquisition conditions has recently been replicated by Kyonka and Grace (2007, 2008).

Grace and McLean (2006) proposed a decision model that could account for response allocation consistent with both generalized matching and categorical discrimination. Their model assumes that response allocation is determined by the relative response strength of the initial-link schedules (i.e., the relative propensity to respond to each alternative). Response strength for a particular initial link is updated after reinforcement has been obtained in the preceding terminal link, depending on the duration of the terminal-link delay. According to the model, subjects make a ''decision'' as to whether the preceding delay was short or long relative to the history of reinforcement delays across both alternatives. If the delay is judged as short, response strength for the associated initial link increases; if the delay is judged as long then response strength decreases. Changes in response strength are made according to a linear-operator rule (with parallel equations for left and right alternatives):

$$RS_{n+1} = RS_n + p_s * (Max_{RS}\text{-}RS_n)$$
$$*\Delta - (1\text{-}p_s) * (RS_n - Min_{RS}) * \Delta \qquad (2)$$

According to Equation 2, RS_{n+1} (expected response strength for cycle n+1) is determined by response strength on the previous cycle (RS_n), modified by an additive (or subtractive) term, depending on whether the delay was judged as short (or long). Specifically, if the previous delay was judged as short (with probability p_s), the response strength increases by a constant fraction (determined by a learning rate parameter, Δ) of the difference between the maximum response strength (Max_{RS}) and current response strength. Conversely, if the previous delay was judged as long (with probability 1-p_s), response strength decreases by a constant fraction of the difference between the current and minimum response strength (Min_{RS}). Whether a delay is classified as short or long depends on a comparison with the distribution of delays experienced across both alternatives. To represent the history of

reinforcement delays, the model uses a log normal distribution with a mean (criterion) equal to the log geometric mean of the experienced delays. The probability that a delay is judged short is the area under the distribution to the right of the delay. The standard deviation of the distribution (σ) is a parameter and determines the accuracy with which delays are classified as short or long.

Grace and McLean (2006) showed that their model could predict response allocation that conformed to generalized matching or categorical discrimination, depending on the value of σ. When σ was relatively low, classification decisions were accurate and response allocation was a nonlinear (sigmoidal) function of the log immediacy ratio. When σ was relatively large, decisions were less accurate, and response allocation was approximately a linear function of the log immediacy ratio. They also showed that the model provided a reasonably good fit to the data from individual subjects.

Christensen and Grace (2008) extended Grace and McLean's (2006) model by proposing that the distribution representing reinforcement history include the intervals between all stimulus transitions. Specifically, they proposed that the criterion against which terminal-link delays were judged as short or long was determined by the delays between initial-link onset and terminal-link entry, in addition to the delays between terminal-link entry and reinforcement. Christensen and Grace showed that with this assumption, the decision model predicted that preference for a constant pair of terminal links was a bitonic function of initial-link duration. Over a substantial range of initial-link durations, preference decreased as initial-link duration increased—the well-known ''initial link effect'' (Fantino, 1969). However, the model predicted a downturn in preference for short initial-link durations, which was confirmed by two experiments.

Although Christensen and Grace's (2008) proposal addresses some of its limitations, the decision model is still inadequate as a general model for concurrent chains. One problem is that the model includes no mechanism for changes in preference across sessions. With the PRBS procedure, relative terminal-link immediacy is not predictable from prior sessions; thus, after sufficient training, response allocation is controlled by the current-session immediacy ratio with little or no detectable effect of history. Consequently, Grace and McLean (2006) assumed that response strength for both alternatives was reset to an intermediate value at the start of each session ($[\text{Max}_{RS} + \text{Min}_{RS}] / 2$). This assumption cannot be valid for steady-state designs, in which the terminal-link schedules remain unchanged for 20 sessions or more. Here the terminal-link immediacy ratio is usually the same as that in the prior session (except for the start of a new condition). To account for the gradual acquisition of steady-state preference (e.g., Grace, 2002), changes in response strength that occur within sessions need to persist, at least to some degree, across sessions.

A simple way to extend Grace and McLean's (2006) model to account for changes in response strength across sessions is to assume that a constant fraction of the change in response strength during a session is retained at the start of the next session. Specifically:

$$RS_{startN+1} = RS_{startN} + (RS_{endN} - RS_{startN}) * \Delta s \tag{3}$$

where RS_{start} and RS_{end} are response strengths at the start and end of the session (subscripted N or $N+1$) respectively, and Δs is a learning rate parameter. With the addition of Equation 3, Grace and McLean's decision model can describe both within- and between-session learning. Note that Δs is assumed to be generally less than 1, so that response strength at the start of session $N+1$ will have regressed back towards the response strength at the start of the previous session. Thus the model predicts spontaneous recovery in choice behavior (Mazur, 1996).

It is important to note that the extended decision model (ExtDM) was developed independently of steady-state models for choice. Unlike LINOP, it is not based on the assumption that terminal-link stimuli acquire conditioned reinforcing value, which in turn is a function of the distribution of reinforcer delays. Instead, the model assumes that the relative likelihood of responding to each initial link is updated according to a series of binary decisions. The purpose of the present research was to compare predictions of the ExtDM and LINOP for a situation that is intermediate between steady-state designs and the PRBS procedure used by Grace et al.
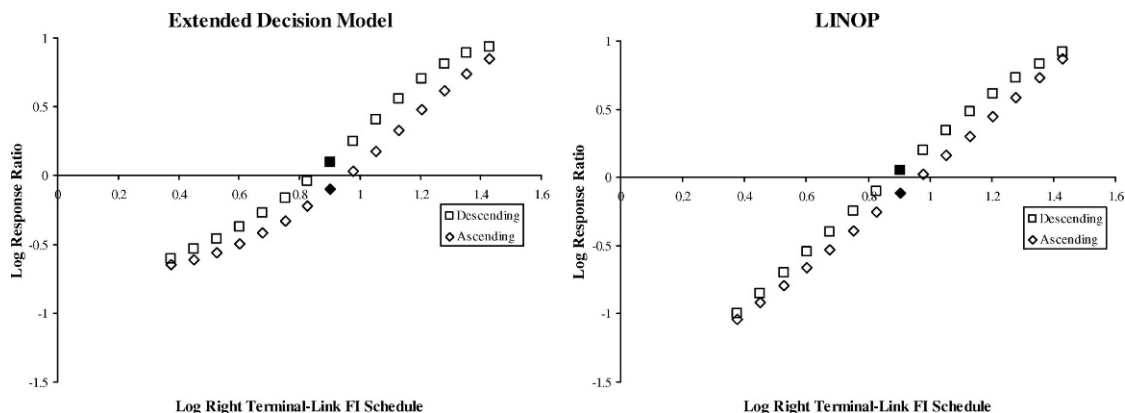
Fig. 1. Log initial-link response ratios as a function of log terminal-link immediacy ratios predicted by ExtDM and LINOP. See text for more details.

(2003) and Grace and McLean (2006). Specifically, we studied changes in response allocation when the relative terminal-link reinforcer immediacies followed a systematic ascending and descending series. This is an intermediate situation because the terminal-link immediacy ratio changes every session, but the changes are correlated because they follow a predictable pattern. In our experiment, the left terminal link was always FI 8 s while the schedule value for the right terminal link changed from 2 s to 32 s and back to 2 s (or from 32 s to 2 s to 32 s) through a geometrically-spaced 17-step series.

Figure 1 shows predictions for this situation by the ExtDM (left panel) and LINOP (right panel). In both panels, the log initial-link response ratio is plotted as a function of the log FI schedule value for the right terminal link. Predictions are shown for the 15 values between 2 s and 32 s which were arranged during both the descending and ascending series. Predictions depend on the specific parameter values used, but the qualitative trends evident in Figure 1 are robust[1]. Both

models predict that preference for the left terminal link (FI 8 s) is overall greater on the descending than ascending series. This would correspond to a hysteresis effect; at the start of the descending series, the right terminal link from the previous session is FI 32 s, and so a nearly maximal preference for the left alternative should have been reached. However, the models differ in terms of the strength of preference for the shorter terminal link depending on whether the right terminal link is less than or greater than 8 s. The filled symbols in Figure 1 indicate when the right terminal link was 8 s, and divide both series into halves in which the absolute values of the log immediacy ratios are equal. According to the LINOP model, the strength of preference for the left alternative when the right delay is greater than 8 s (i.e., points to the right of the filled symbols) is the same as the strength of preference for the right alternative when the right delay is less than 8 s (i.e., points to the left of the filled symbols). However, the ExtDM predicts that the strength of preference should be overall greater when the right delay is longer than 8 s. This exemplifies the terminal-link effect (MacEwen, 1972; Grace, 2004; Grace & Bragason, 2004)—that preference should be more extreme with overall longer delays, with the delay ratio held constant.

[1] Parameter values for the models were as follows: For the ExtDM, $\sigma = 0.3$, $\Delta_S = 0.3$, $\Delta = 0.3$, and the maximum and minimum response strengths for both alternatives were 1.0 and 0.05, respectively. For LINOP, value was defined as a power function of reinforcer immediacy with exponent = 2, and the learning rate parameters Δ and Δ_S were 0.5 and 0.3, respectively. To simulate each session, the models' predictions were computed over 12 cycles, corresponding to 12 blocks of 6 cycles in concurrent chains. The predicted log response ratio was calculated for each cycle, and then averaged across cycles to give a value for the session.

METHOD

Subjects

Eight pigeons of mixed breed, numbered 221, 222, 223, 224, and 191, 192, 193, 194,

served as subjects and were maintained at 85% of their free-feeding weight +/- 15 g through appropriate postsession feeding. Subjects were housed individually in a vivarium with a 12h:12h light/dark cycle (lights on at 0600), with water and grit freely available in the home cages. Pigeons 221, 222, 223, and 224 (Group Experienced) were experienced with rapid-acquisition concurrent-chains procedures and had served as subjects in Grace, Bragason, and McLean's (2003) research just prior to the start of the present study, whereas Pigeons 191, 192, 193 and 194 (Group Naïve), although experienced with other procedures, had no prior training with rapid-acquisition concurrent chains.

Apparatus

Four standard three-key operant chambers, 32 cm deep × 34 cm wide × 34 cm high, were used. The keys were 21 cm above the floor and arranged in a row. In each chamber there was a houselight located above the center key that provided general illumination, and a grain magazine with an aperture centered 6 cm above the floor. The magazine was illuminated when wheat was made available. A force of approximately 0.15 N was necessary to operate each key. Each chamber was enclosed in a sound-attenuating box, and ventilation and white noise were provided by an attached fan. Experimental events were controlled and data recorded through a microcomputer and MEDPC® interface located in an adjacent room.

Procedure

For all pigeons, training started immediately in a concurrent-chains procedure. The houselight provided general illumination at all times except during reinforcer delivery. With few exceptions, sessions were run daily and at approximately the same time (1000h for Group Experienced; 1200h for Group Naïve).

Sessions ended after 72 initial- and terminal-link cycles or 70 min, whichever occurred first. At the start of a cycle, the side keys were illuminated white to signal the initial links. An entry was assigned pseudorandomly to the left or right terminal link with the constraint that in every six cycles, three entries occurred to each terminal link. An initial-link response produced an entry into a terminal link

provided that: (a) it was made to the preselected key; (b) an interval selected from a VI 10-s schedule had timed out; and (c) a 1-s changeover delay (COD) was satisfied— i.e., at least 1 s had elapsed following a changeover to the side for which terminal-link entry was arranged.

The VI 10-s initial-link schedule did not begin timing until the first response had occurred in each cycle, to allow any pausing after the completion of the previous terminal link to be excluded from initial-link time. The VI 10-s schedule contained 12 intervals constructed from an exponential progression (Fleshler & Hoffman, 1962). Separate lists of intervals were maintained for cycles in which the left or right terminal link had been selected, and were sampled without replacement so that all 12 intervals would be used three times for both the left and right terminal links each session.

When a terminal link was entered, the color of the side key was changed (left key to red, right key to green) while the other key was darkened. Terminal-link responses were reinforced according to FI schedules. When a response was reinforced all lights in the chamber were extinguished, and the grain magazine raised and illuminated for 3 s.

The FI schedule value for the red (left) terminal link was always 8 s, and the value for the green (right) terminal link was one of the following: 2, 2.38, 2.83, 3.36, 4, 4.76, 5.66, 6.73, 8, 9.51, 11.31, 13.45, 16, 19.03, 22.63, 26.91, or 32 s. The right terminal-link schedule values were equally spaced in logarithmic terms, and occurred in an ascending or descending series across sessions. For example, the 2-s delay was always followed by 2.38 s in the next session, and 2.83 s in the session after that (i.e., in the order listed above), whereas the 32-s delay was always followed by delays in the reverse order (i.e., 26.91 s in the next session, 22.63 s in the session after that, etc.).

For 2 pigeons in Group Experienced (221 and 222), the right terminal link began at 2 s and three ascending/descending series were completed; for Pigeons 223 and 224, the right terminal link began at 32 s and three descending/ascending series were completed. All pigeons in Group Naïve first received 21 sessions in which both terminal-link schedules were FI 8 s. The purpose of this training was to establish a baseline from which the effects of

the ascending or descending series could be assessed. Delays were then increased across sessions to 32 s for Pigeons 191 and 192 according to the geometric series, and decreased across sessions to 2 s for Pigeons 193 and 194. All pigeons then completed three descending/ascending series (191 and 192) or ascending/descending series (193 and 194).

RESULTS

Figure 2 shows response allocation and the programmed immediacy ratio plotted over sessions for all subjects across the three ascending and descending series. Figure 2 illustrates that response allocation for all subjects in both groups tracked changes in the immediacy ratio. Response allocation increasingly favored the left initial link during the ascending series (in which the right terminal link changed from 2 s to 32 s), and the right initial link during the descending series (in which the right terminal link changed from 32 s to 2 s). Individual differences are also apparent. For example, shifts in response allocation were small and gradual across sessions for some pigeons (e.g., 222, 223, 224 in Group Experienced, and 192 in Group Naïve), corresponding to changes in the log immediacy ratio, but large changes were evident for others (e.g., Pigeons 221 and 193). There was a pronounced bias toward the left initial link for Pigeon 223, and to a lesser extent for Pigeons 193 and 194. Overall, there appears to be no systematic difference between Group Experienced and Naïve in terms of changes in response allocation across sessions.

To assess results in Figure 2 more systematically, individual-subject data were entered into a repeated-measures analysis of variance (ANOVA) with group (Naïve or Experienced) as a between-subjects factor and log immediacy ratio, replication (first, second, or third presentation of a series) and series (ascending or descending) as within-subjects factors. The main effects of series and log immediacy ratio were significant, $F(1,6) = 28.96$ and $F(14,84) = 42.73$, both $p < 0.01$, respectively, whereas those of group and replication were not, $F(1,6) = 1.23$ and $F(2,12) = 3.02$, both *ns.*

There were two significant interactions, replication x log immediacy ratio, $F(28,168) = 1.55$, $p < 0.05$, and series x log immediacy ratio, $F(14,84) = 4.39$, $p < 0.01$. Analysis of simple effects showed that response allocation favored the right initial link relatively more during the second replication when the delay was 8 s and 9.51 s, and favored the left initial link relatively more during the third replication when the delay was 16 s and 22.63 s. Although reasons for these differences are unclear, the effects were small and apparently unsystematic in the context of the overall changes in preference.

To highlight the series x log immediacy interaction, Figure 3 shows log response ratio as a function of the log terminal-link immediacy ratio, averaged across replications. All subjects responded relatively more to the left initial link to a greater extent during the descending series, especially for immediacy ratios in the middle of the range, but response allocation tended to converge at the most extreme immediacy ratios. Overall, the pattern might be described as a ''bubble'' near the middle of the immediacy ratio range, and indicates a hysteresis effect. This effect occurred as follows: At the end of the ascending series, the right terminal-link delay was 32 s and response allocation strongly favored the left key. The preference for the left key persisted while the right-key delay decreased during the descending series, but eventually responding switched to favor the right when the delay became sufficiently short. When the delay was 2 s at the end of the descending series, response allocation strongly favored the right key. As the delay began to increase in the ascending series, preference for the right key persisted until the delay became sufficiently long, when it switched to the left key. Thus, the persistence in response allocation at the end of both series produced an overall increased preference for the left key in the descending series, creating the bubble pattern. The magnitude of this effect varied across subjects; it was strong for Pigeons 221 and 191, but relatively weak for Pigeons 222 and 223. Nevertheless, results for all subjects showed evidence of hysteresis.

To quantify the magnitude of the hysteresis effect, we calculated the delay associated with the midpoint of the range in response allocation for both ascending and descending series (averaged across replications). Specifically, we computed the average of the log response ratios for the two most extreme delays in each series and then, using linear
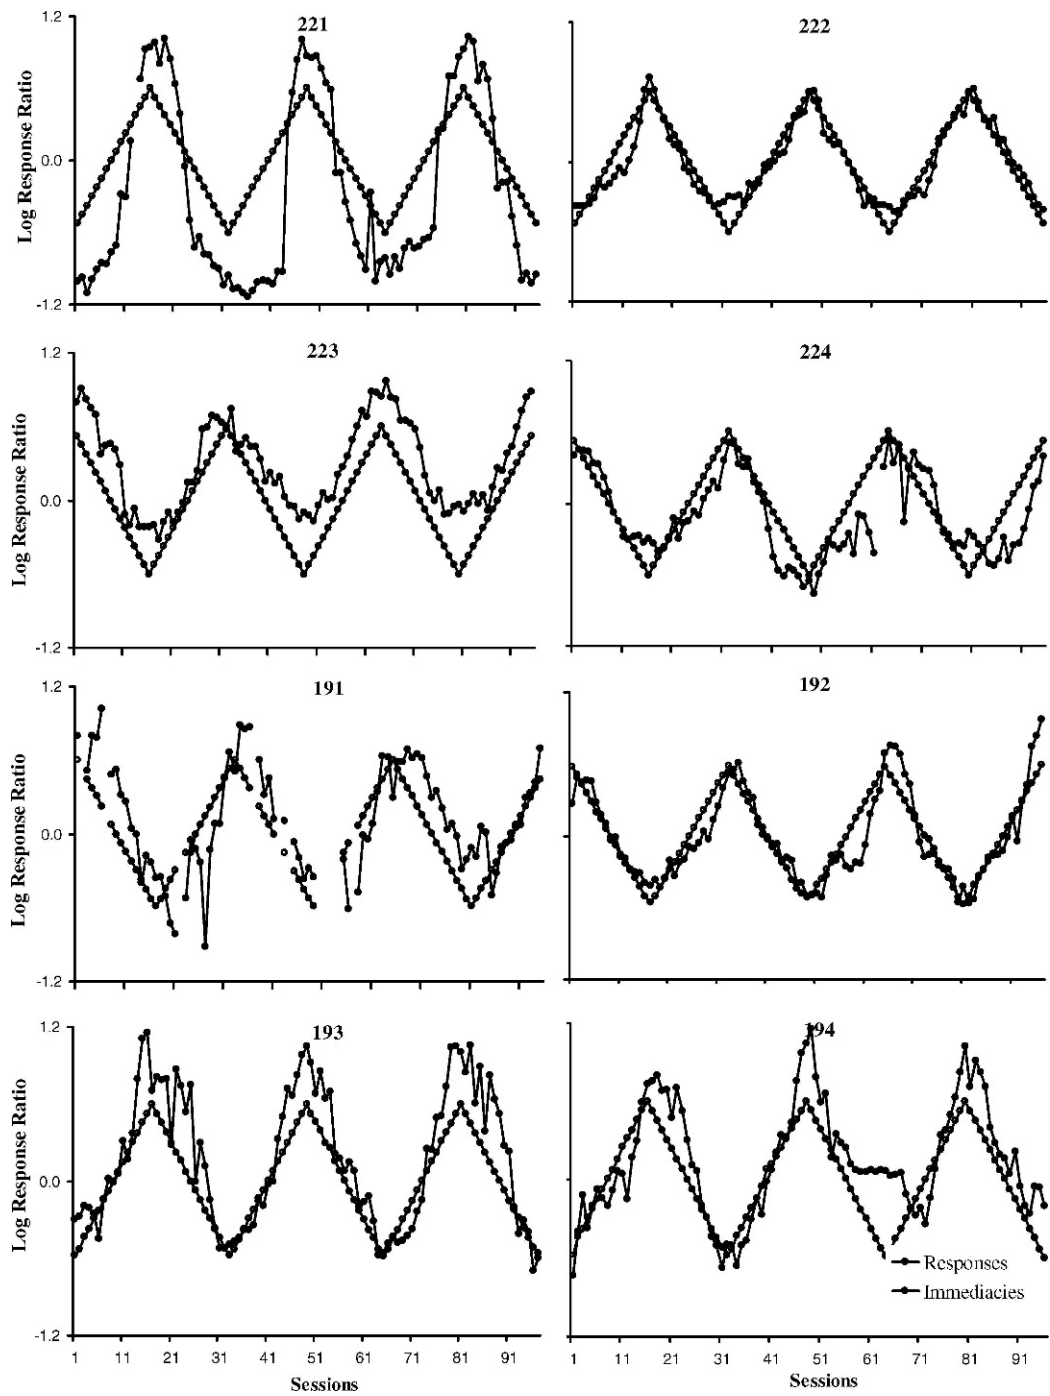
Fig. 2.   Obtained log initial-link response allocation and log programmed terminal-link immediacy ratios across all three replications of the ascending/descending series for subjects in Group Experienced (Pigeons 221, 222, 223, and 224) and Naïve (Pigeons 191, 192, 193, and 194).
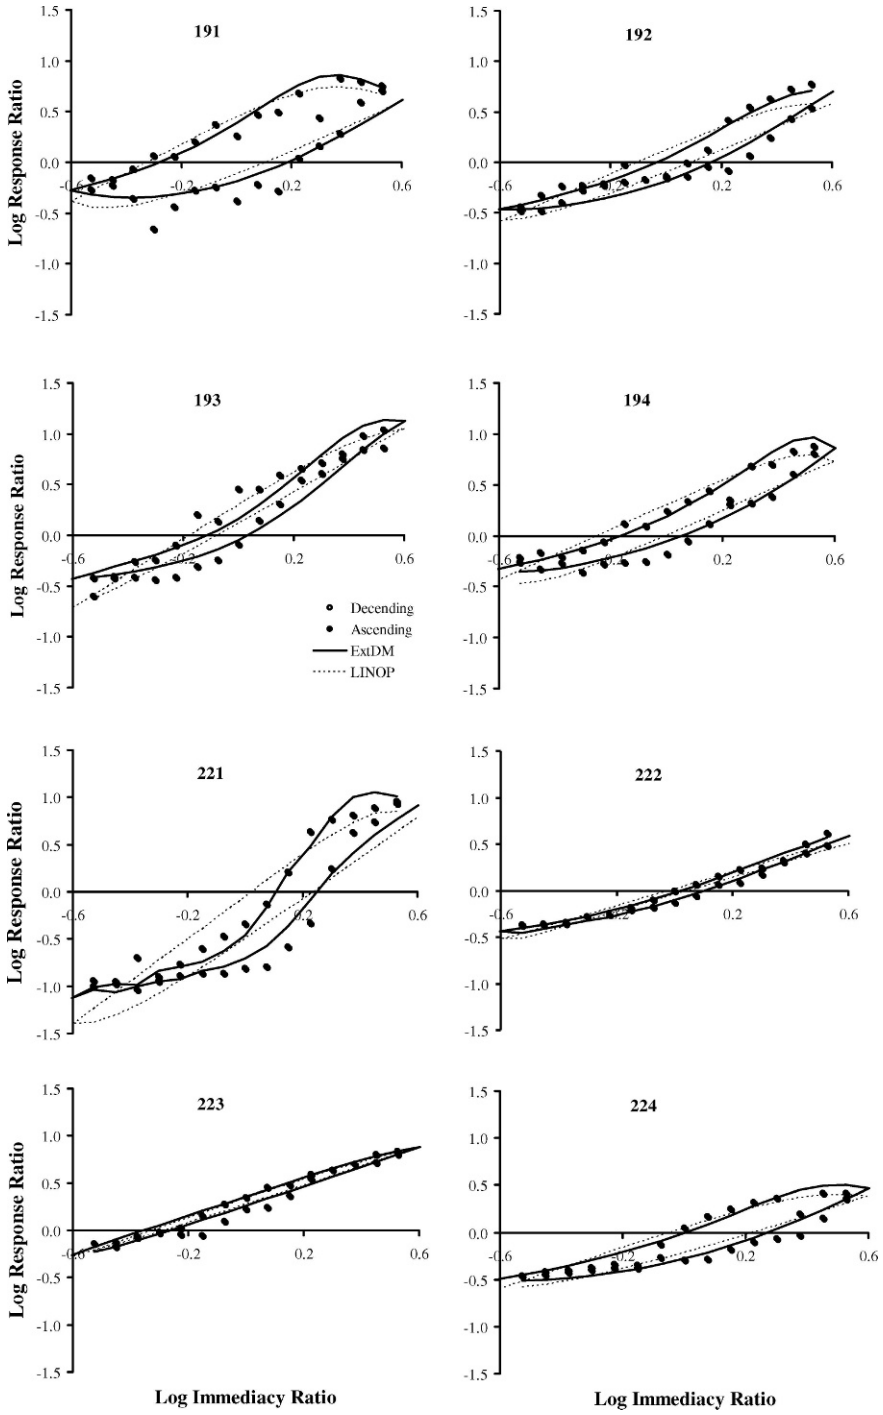
Fig. 3. Obtained log initial-link response ratios as a function of programmed log terminal-link immediacy ratios, for both ascending and descending series, averaged across replications for individual subjects. Predictions of ExtDM and LINOP are also shown by solid and dashed lines, respectively.

Table 1

Midpoint terminal-link delays (in seconds) and corresponding log response ratios for ascending and descending series, for all subjects.

	Delay			Log Response Ratio	
Pigeon	Ascending	Descending	Pigeon	Ascending	Descending
221	14.92	10.09	221	0.24	−0.13
222	15.11	10.65	222	0.17	0.07
223	10.95	6.85	223	0.37	0.28
224	21.11	7.97	224	0.15	−0.13
191	14.79	5.44	191	0.17	0.06
192	17.21	11.47	192	0.24	0.12
193	10.62	5.03	193	0.31	−0.09
194	18.42	9.95	194	0.39	0.33
Mean	15.39	8.43	Mean	0.26	0.06
SE	1.26	0.87	SE	0.03	0.06

interpolation, found the delay that corresponded to this log ratio.

The midpoint delays and the corresponding log response ratios are listed in Table 1. For all subjects, midpoint delays were greater for the ascending ($M = 15.39$, $SE = 1.26$) compared to the descending series ($M = 8.43$, $SE = 0.87$). A repeated-measures ANOVA with group as a between-subjects factor and series as a within-subjects factor found a significant effect of series, $F(1,6) = 34.37$, $p < 0.01$, but the effect of group and the series x group interaction were nonsignificant. This demonstrates that the magnitude of the hysteresis effect was substantial, encompassing approximately one-quarter of the range of variation in log immediacy ratio.

Next we compared the ability of LINOP and the ExtDM to provide a quantitative account of the present data. Log initial-link response ratios were computed for every block of six cycles in each session (i.e., 12 blocks per session), then averaged across replications, giving a total of 384 data points (12×32) for each subject. The models were then fitted by obtaining parameter estimates that maximized the variance accounted for in the data. For LINOP the parameters included Δ_S, which determined the rate of learning across sessions, and Δ_R, which determined the rate of learning within sessions. There was also a sensitivity exponent, q, in the function determining the asymptotic value of a delayed reinforcer, $V = 1 / (c + d^q)$, where delay is d seconds and c is an additive constant which was set equal to 0 for the fits presented here (see Grace, 2002, Equation 4). An additive bias parameter, log b, was also used. Parameter estimates that maximized the variance accounted for were obtained through nonlinear optimization (Microsoft Excel© Solver).

Figure 4 shows the obtained log initial-link response ratios as a function of LINOP predictions for the individual block data (i.e., session twelfths). Slopes for best-fitting regressions are also shown, and the slopes are close to 1.0, suggesting that the LINOP model captured the overall trends in the data. Averaged across subjects, the LINOP model accounted for 91% of the variance in the session-12[th] data. However, there is some evidence of sigmoidal curvature for some subjects in Figure 4, indicating that the LINOP predictions deviate systematically from the obtained data. For example, Pigeons 221, 224, 192 and 194 appear to have obtained data that follow a trend that begins below the regression line at low predicted values and as the predictions increase, falls above the regression line. These subjects also have the most pronounced bubble between series in the session data (Figure 3), and suggest that LINOP struggles to describe hysteresis effects in response allocations. Parameter values for the fits of the LINOP model to the session-12th data are listed in Table 2.

The dashed lines in Figure 3 show the whole-session average values (obtained and predicted by the LINOP model) as a function of the log immediacy ratio for both the ascending and descending series. Predicted values were calculated by averaging across the predicted values for the session-12th data. Overall, LINOP provided a reasonably good account of the data, accounting for 85% of the variance. LINOP was able to predict the
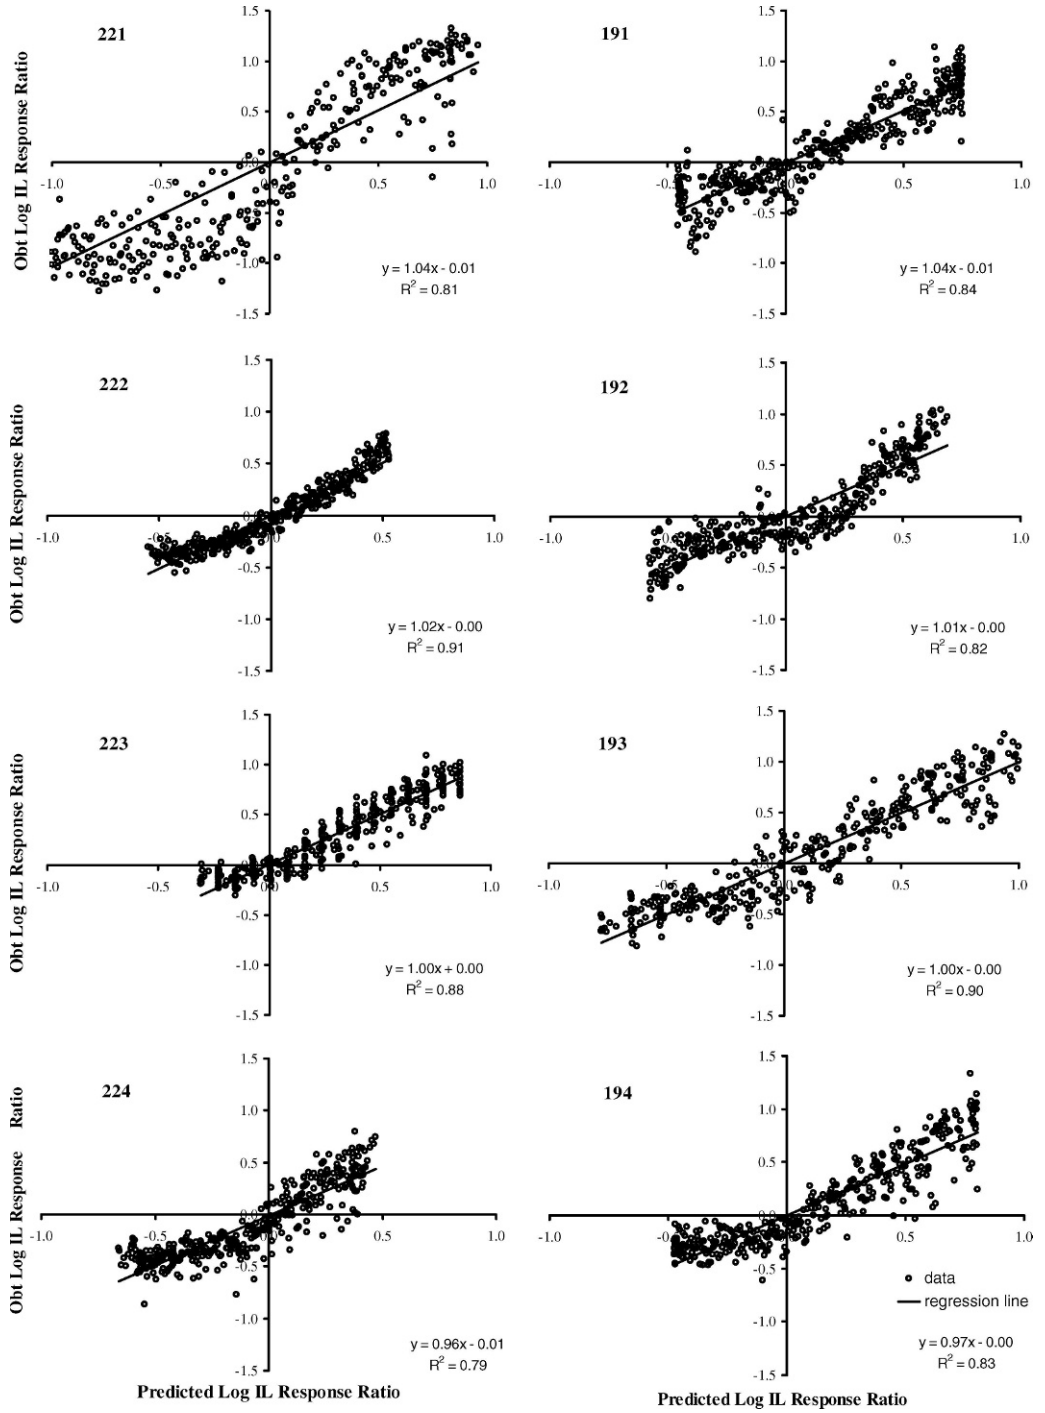
Fig. 4.   The session-12th obtained log initial-link response ratios as a function of LINOP-predicted log immediacy ratios for both ascending and descending series, averaged across replications for individual subjects. Included are the regression lines, associated best fitting r^2 values and linear regression parameters.

Table 2

LINOP parameter and VAC values for fits to session 12[th] data.

	LINOP					
Session						Session 12th
VAC	q	log b	Δ	Δs		R^2
Pigeon						
221	0.89	2.25	−0.19	0.13	0.54	0.81
222	0.95	0.91	−0.01	0.17	1.00	0.91
223	0.96	1.04	0.32	0.58	0.27	0.88
224	0.89	1.09	−0.05	0.13	0.36	0.79
191	0.87	1.41	0.23	0.03	1.00	0.84
192	0.89	1.18	0.04	0.17	0.38	0.82
193	0.95	1.70	0.25	0.37	0.28	0.90
194	0.92	1.23	0.18	0.05	1.00	0.83
Average	0.92					0.85

separation between ascending and descending series in the full-session data, corresponding to the hysteresis effect. Additionally, LINOP was able to capture some of the nonlinearity in the full-session data (e.g., see Pigeons 221, 224, and 194). However, LINOP appears to fail in capturing some of the patterns of hysteresis. In particular, the response patterns of Pigeons 221, 224, 191, 193 and 194 appear to show evidence of little change in response ratios in the beginning of the ascending series. LINOP seems only able to capture this effect for Pigeons 191 and 194 and fails to describe the hysteresis in the ascending series for Pigeons 221, 224, and 193.

Next we applied the extended version of Grace and McLean's (2006) decision model (Equations 1–2) to the data. The criterion value was calculated for each session as the average of the log programmed intervals between stimulus transitions (i.e., initial-link onset to terminal-link entry, and terminal-link entry to reinforcement). The probability of the relative current delay being judged short relative to the criterion was then used in the prediction of preference for the session-12th data. The maximum and minimum response strengths for both alternatives were initially set equal to 1.0 and 0.01, respectively. Solver was used to obtain best-fitting values of the standard deviation (σ), learning rate parameter for the terminal links (Δ) and between-session changes (Δ_S), as well as an additive bias parameter (log b). Thus, both the ExtDM and LINOP had four free parameters. Parameter values for the fits to the individual data are listed in Table 3.

Figure 5 shows obtained log initial-link response ratios as a function of ExtDM predictions for the session-12th data. The best-fitting regression lines are also shown. Overall, the ExtDM did a good job of describing the session-12th data, accounting for an average of 88% of the variance. The regression slopes were also all close to 1.0. However, there is some evidence of curvature in the scatterplots that indicate that predictions of the ExtDM, like those for LINOP, sometimes deviate systematically from the obtained values. For example, Pigeons 192, 193 and 223 appear to have obtained data that follow a trend that begins below the regression line at low predicted values and as the predictions increase fall above the regression line.

The solid lines in Figure 3 show the resulting session average values as a function of the log immediacy ratio. The ExtDM provided an excellent account of the data, with an average VAC of 95% for the full-session data. The ExtDM provided a good description of results for subjects for which there was a clear separation between the ascending and descending series, as well as when the series nearly superposed. For example, Pigeon 224 has a distinct separation between series, while Pigeon 223 has almost identical curves for the ascending and descending series. Compared to the LINOP predictions Figure 3 appears to show the ExtDM is able to capture both patterns of responding. Moreover, the ExtDM also appears to capture hysteresis in both ascending and descending series. This is most evident in subject 221, whose obtained and predicted curves become flatter at the start of

Table 3

ExtDM parameter and VAC values for fits to session 12[th] data.

	ExtDM						
	Session						Session 12th
	VAC	Log C	σ	log b	Δ	Δs	R^2
Pigeon							
221	0.97	0.90	0.09	−0.64	0.21	0.41	0.90
222	0.99	0.90	0.34	−0.06	0.10	1.49	0.95
223	0.97	0.90	0.34	0.31	0.27	0.45	0.88
224	0.93	0.90	0.26	−0.16	0.11	0.36	0.83
191	0.89	0.90	0.17	0.04	0.03	0.69	0.85
192	0.95	0.90	0.22	−0.09	0.18	0.26	0.88
193	0.92	0.90	0.16	0.04	0.26	0.22	0.88
194	0.96	0.90	0.25	0.06	0.04	1.50	0.87
Average	0.95						0.88

both ascending and descending series. In addition, the ExtDM seems to be a good approximation of more linear patterns of response allocation, for example with Pigeons 222 and 223.

Comparing the model fits, those for the ExtDM were overall superior, with higher VAC for 7 of 8 pigeons for the full-session data, and for 6 of 8 (with one tie) for the session-12th data. Because the models have the same number of free parameters, this suggests that the ExtDM may provide a better description of response allocation for the present data.

However, even if two models have the same number of parameters, one may have greater flexibility in terms of being able to predict a greater range of outcomes (Pitt, Myung, & Zhang, 2002). If so, the model may account for a higher percentage of variance than a competitor because of this flexibility. Thus, to determine whether the ExtDM and LINOP differed in terms of flexibility, we conducted an analysis in which both models were fitted to simulated data generated by each model. The simulated data were obtained by adding random noise (distributed uniformly between −0.1 and 0.1) to the predicted values when each model was fitted to the average session-12th data. If either model is more flexible, then it should provide not only the best account of simulated data generated by that model, but an equal or better account of data generated by the other model as well. For simulated data generated from the ExtDM, the ExtDM and LINOP accounted for 93.5% and 86.4% of the variance, respectively. For simulated data generated from LINOP, the ExtDM and LINOP

accounted for 87.1% and 93.5% of the variance. In both cases, the model that generated the simulated data provided the better fit. This suggests that there is no difference in flexibility between ExtDM and LINOP. We therefore conclude that the ExtDM provides a better overall account of the present data.

Finally we examined whether the terminal-link effect—that is, a stronger preference for the shorter terminal-link delay when the absolute values of the delays increase with their ratio held constant—was obtained in the present data and whether the models could account for the result. Figure 6 shows the obtained log response ratios (full session) as a function of the absolute value of the log immediacy ratio for individual subjects, separately, according to whether the terminal-link FI schedule for the right alternative was less than or greater than 8 s. Each data point represents an average across the ascending and descending series. Because the 8-s duration was the midpoint of both series, the log immediacy ratios formed pairs with equal absolute values. The terminal-link effect predicts that sensitivity to the log immediacy ratio, as measured by the slope of the generalized-matching function of the log immediacy and log response allocation, should be greater when the right terminal-link schedule was greater than 8 s compared to when it was less than 8 s. Figure 6 shows that for all subjects the >8 s log response ratios had a greater slope than the corresponding < 8 s response ratios. Thus the data exemplified the terminal-link effect, that is, preference was more extreme with longer absolute terminal-link duration as relative duration was held constant.
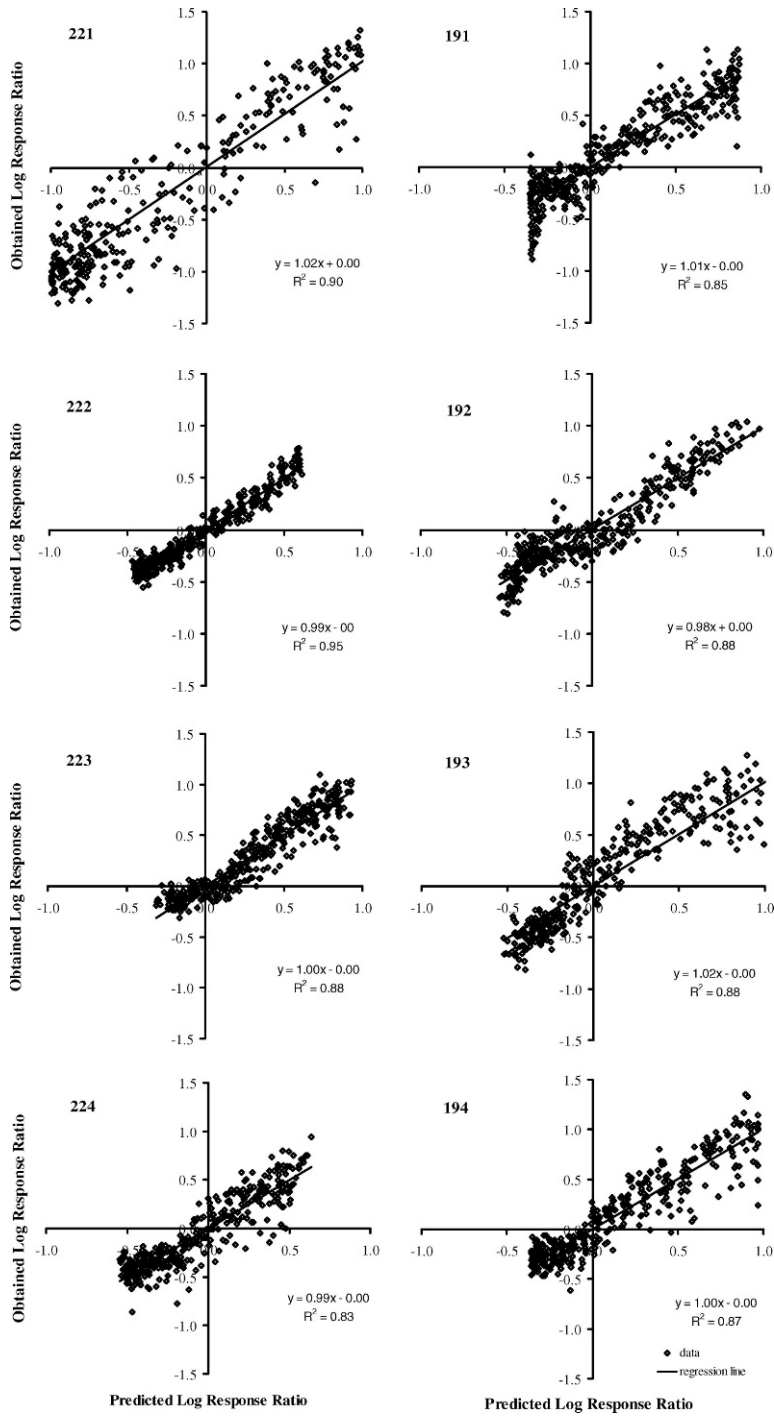
Fig. 5.   The session-12th obtained log initial-link response ratios as a function of ExtDM-predicted log immediacy ratios for both ascending and descending series, averaged across replications for individual subjects. Included are the regression lines, associated best fitting r^2 values and linear regression parameters.
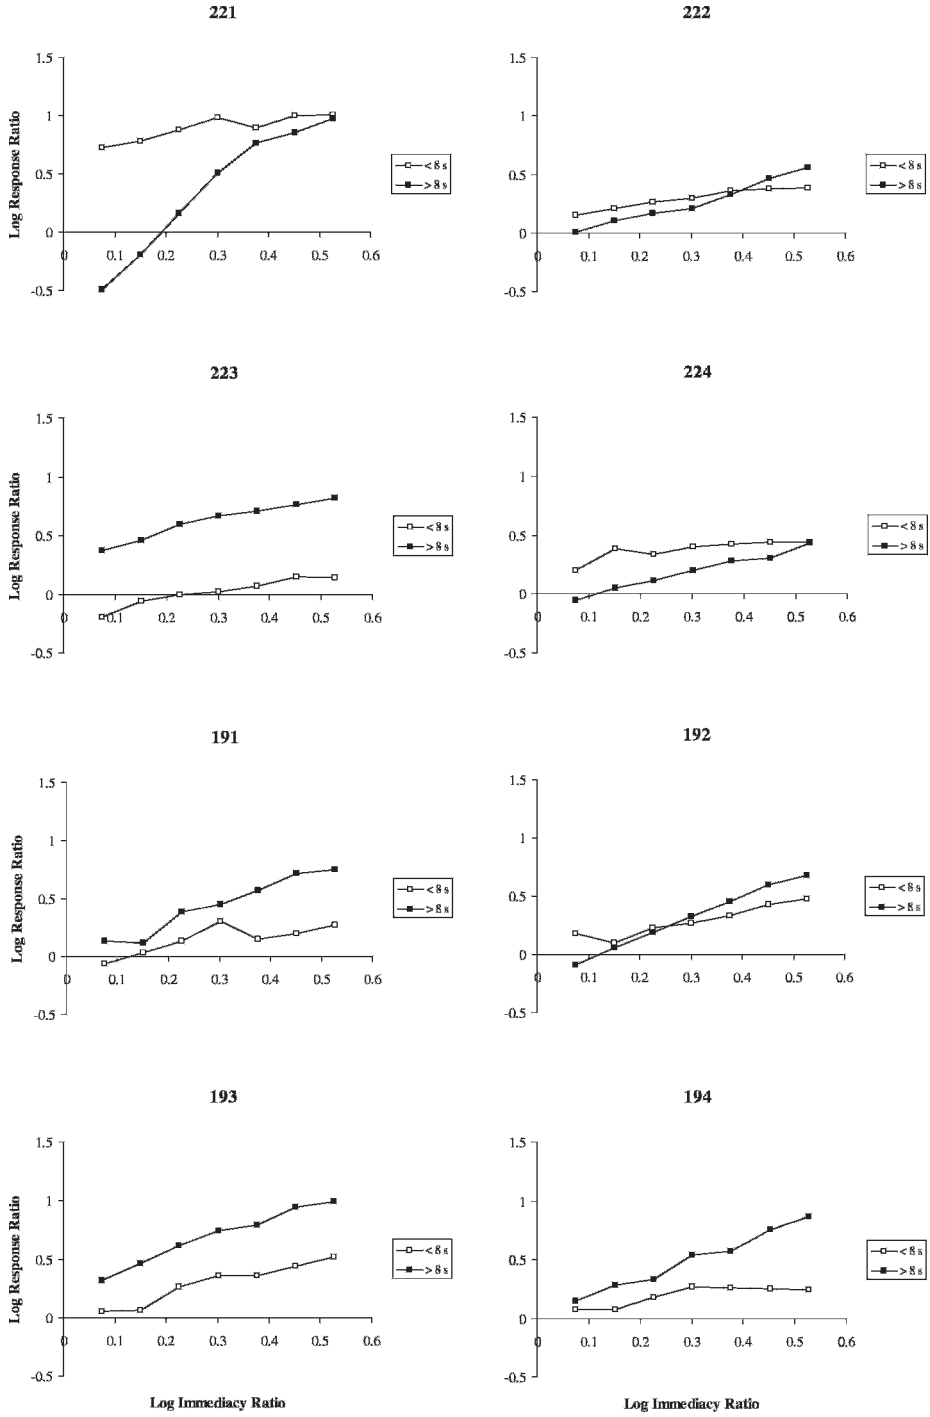
Fig. 6. Log initial-link response ratios as a function of log terminal-link immediacy ratios for which the right terminal link FI schedule was greater than or less than 8 s, for both ascending and descending series, averaged across replications for individual subjects.
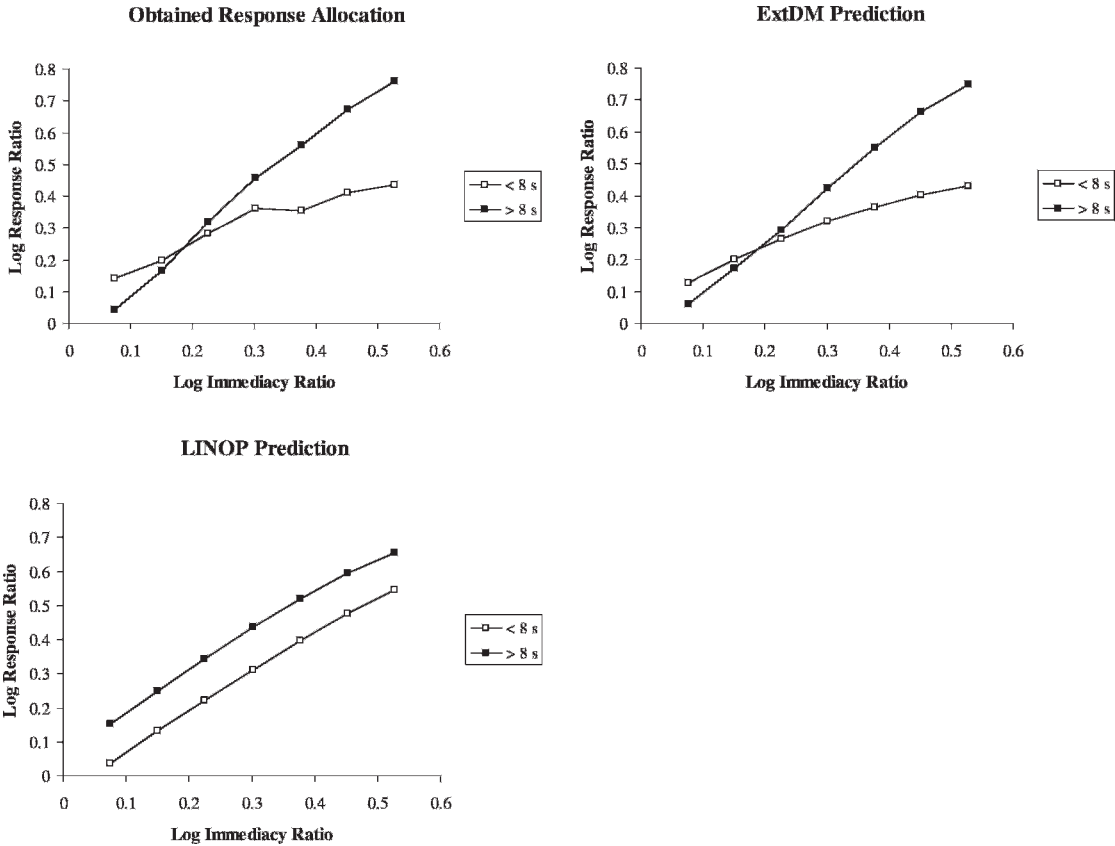
Fig. 7. Obtained log initial-link response ratios (top panel) and predictions of ExtDM (upper right panel) and LINOP (bottom panel) as a function of log terminal-link immediacy ratios for which the right terminal-link FI schedule was greater than or less than 8 s, for both ascending and descending series, averaged across replications and subjects.

Figure 7 shows the group averages for predicted and obtained log response ratios for both series when the terminal-link duration was less than or greater than 8 s (upper left panel). Like the individual data, the obtained data show steeper slopes in the >8 s than the <8 s log response ratios. The ExtDM predictions in the upper right panel (obtained by averaging across predictions for the ascending and descending series in Figure 3) show the same pattern as the obtained data. However, the corresponding LINOP predictions have parallel slopes for the two sets of conditions, indicating that LINOP failed to predict the terminal-link effect.

DISCUSSION

The present study explored how initial link response allocation in concurrent chains changed when one terminal-link delay followed an ascending and descending sequence across sessions while the other remained constant. Our goal was to test predictions of two models for acquisition in concurrent chains: an extended version of the decision model proposed by Grace and McLean (2006) and Christensen and Grace (2008), and the LINOP model (Grace, 2002). The decision model had previously been applied only to situations in which the terminal links changed unpredictably across sessions. Here, we assumed that a proportion of the change in response strength within a session would be retained at the start of the next session.

The terminal-link schedule for the left alternative was always FI 8 s, while the right terminal-link schedule varied between FI 2 s and FI 32 s according to a geometric series. Two predictions of the ExtDM were evaluated: that a

hysteresis or carryover effect would be obtained, and that response allocation would be more sensitive to changes in the immediacy ratio at higher absolute terminal-link durations (see Figure 1). Both predictions were confirmed.

For all subjects, scatterplots of the log initial-link response and log immediacy ratios showed a gap or ''bubble'' between data for the ascending and descending series (see Figure 3). This phenomenon occurred because the series tended to converge at the extreme immediacy ratios, whereas for inter-mediate ratios the log response ratio tended to favor the left initial link to a greater extent during the descending series. Because the descending series began after the right-key delay expected to produce maximal prefer-ence for the left key (32 s), the left-key bias during the descending series represents a hysteresis effect. This effect was also exempli-fied by indifference points (i.e., the right terminal-link delay associated with the mid-point of the total shift in preference; see Table 1) that were greater for the ascending than descending series. Both LINOP and the ExtDM predicted the hysteresis effect.

This result is similar to that reported by Field, Tonneau, Ahearn and Hineline (1996), who studied pigeons' choices between FR 30 and VR 60 terminal links in concurrent chains. Across successive phases of their experiment, the VR distribution was manipulated such that the minimum response requirement was changed according to an ascending and descending series. Preference for the VR alternative tracked the minimum requirement; a requirement of 1 produced a strong preference for the VR terminal link, and this preference decreased as the requirement was increased up to 15. Field et al. found that for a given minimum require-ment, preference for the VR alternative was greater on the ascending than descending series, which is analogous to the hysteresis effect reported here. However, one difference is that each phase in Field et al.'s experiment lasted for 11 sessions. Thus, despite the differences in procedure (e.g., interval versus ratio schedules; schedules changed after 1 and 11 sessions), both experiments produced similar hysteresis effects. It is unknown whether such hysteresis depends on how frequently the terminal links are changed.

Overall, the ExtDM provided a very good account of the data in quantitative terms, accounting for an average of 88% of the variance in the session-12th data and 95% of the variance in the session data. These are somewhat higher than the corresponding values for LINOP (85% and 92%), as well as for the fits of the original version of the decision model to Grace and McLean's (2006) data (73% and 84%). However, it is worth noting that there was some evidence of systematic deviation in the obtained versus predicted scatterplots for the session-12th data (see Figure 5), indicating that the ExtDM was unable to capture all of the trends in the data.

We also tested whether preference for the shorter terminal link would increase as overall terminal-link duration increased, with the immediacy ratio held constant. This result is known as the terminal-link effect, and has been one of the most widely-studied phenom-ena in concurrent chains, having been ob-tained when terminal links differ in terms of reinforcer magnitude (Navarick & Fantino, 1976) and probability (Spetch & Dunn, 1989), as well as immediacy (Grace & Bragason, 2004; Grace, 2004; MacEwen, 1972; Williams & Fantino, 1978). In the present experiment, the delays were geometrically spaced so the ratios between 1:1 and 4:1 were the reverse of those between 1:4 and 1:1. Thus we could compare sessions in which the delays were either both less than 8 s, or both greater than 8 s, with the ratio of delays held constant. For all subjects, the slope relating log response allocation to the log immediacy ratio was steeper when the delays were greater than 8 s (Figure 6). This result was predicted by the ExtDM, but not LINOP (Figure 7).

How the ExtDM Accounts for the Terminal-Link Effect

The terminal-link effect has been considered one of the most theoretically interesting results in concurrent chains, because it represents a striking violation of Weber's law in the temporal domain: Relative discrimination (i.e., response allocation) is not constant at constant delay ratios (Gibbon, 1977). No single explanation for the terminal-link effect is universally accept-ed, although it is predicted by all viable models for steady-state choice in concurrent chains such as delay-reduction theory (Fantino, 1969; Fantino & Romanowich, 2007), the contextual choice model (Grace, 1994), and the hyperbolic value-added model (Mazur, 2001). It is thus
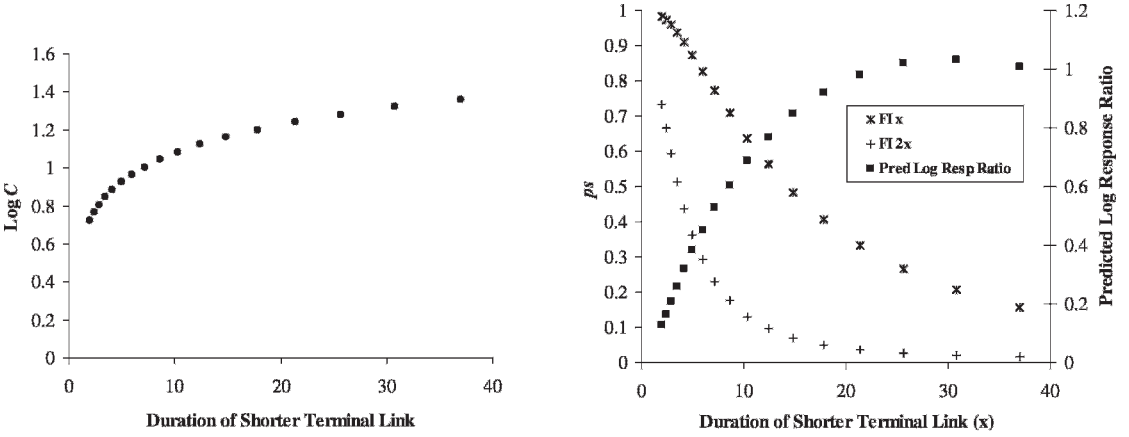
Fig. 8.   Illustration of how the ExtDM predicts the terminal-link effect. The left panel shows Log *C* as a function of the shorter terminal-link delay. The right panel shows the probabilities that the shorter (FI *x*) and longer (FI 2*x*) delays were judged short relative to the criterion (*x*'s and *'s, respectively; left axis), and the resulting predicted log response allocation (filled squares, right axis).

worth considering how the ExtDM is able to account for the terminal-link effect.

Christensen and Grace (2008) showed that the asymptotic response allocation predicted by the ExtDM could be described with the following equation:

$$\frac{B_L}{B_R} = \frac{RS_{asympL}}{RS_{asympR}}$$
$$= \frac{p_{sL}Max_{RS} + (1 - p_{sL})Min_{RS}}{p_{sR}Max_{RS} + (1 - p_{sR})Min_{RS}}, \quad (4)$$

In which B indicates initial-link responses, RS_{asymp} the asymptotic response strength, and p_s is the probability that a terminal-link delay is judged short, subscripted for the left and right alternatives. Max_{RS} and Min_{RS} are the maximum and minimum response strengths (set equal to 1.0 and 0.01, respectively). Equation 4 predicts that response allocation is determined by the relative strength of responding to the initial links, which in turn is calculated as a weighted average of the maximum and minimum response strengths, depending on the probability that a terminal-link delay is judged short relative to the criterion. This probability is the inverse of the cumulative normal distribution with mean equal to the criterion (log C), and a standard deviation σ :

$$ps = 1 - \Phi(\log D, \log C, \sigma) \quad (5)$$

where Φ is the cumulative normal distribution and D is the terminal-link delay to reinforcement.

To illustrate how the ExtDM predicts the terminal-link effect, we used Equations 4 and 5 to calculate the predicted response allocation for a series of terminal-link schedules in which the delay ratio was always 1:2 while the absolute durations varied from FI 2 s FI 4 s to FI 37 s FI 74 s (specific values were determined by a geometric series in which the schedules were increased by 20% at each step). The initial-link schedule was VI 10 s, and $\sigma = 0.2$.

Figure 8 shows how the ExtDM predicts the terminal-link effect. Displayed are the probabilities that reinforcer delays associated with the FI *x* (left) and FI 2*x* (right) terminal links are judged short (p_s) as *x* ranges from 2 s to 37 s. Preference for the shorter terminal link increases as a negatively-accelerated function of *x*, as illustrated by the filled squares in the right panel. Depending on the range of terminal-link durations and specific parameter values chosen, the model can predict a downturn in preference at higher overall durations (i.e., a bitonic function). With VI terminal links, Grace (2004) found that preference increased as a negatively accelerated function of terminal-link duration, and Gentry and Marr (1980) reported that with FI terminal links, preference for some subjects showed a downturn at high absolute durations (although the results were not obtained for all subjects).

The left panel of Figure 8 shows that the criterion (log C) increases monotonically as a function of *x*. The probabilities that delays are

judged short (p_s) are displayed in the right panel. For very small values of x, p_s is high for both schedules because both delays are short relative to log C, which is determined jointly by the initial and terminal-link schedules. As x increases, p_s falls more steeply for the longer terminal link at first, leading to an increasing preference for FI x. However, as durations increase, p_s decreases less rapidly for FI x, resulting in a flattening and eventual downturn in preference. Thus, the model predicts that the shape of the terminal-link effect arises from different rates of change in p_s for the two alternatives as the overall terminal-link duration increases.

The explanation for the terminal-link effect provided by the decision model resembles, to some extent, that provided by delay-reduction theory (DRT; Fantino, 1969; Fantino & Romanowich, 2007). Log C plays a role similar to that of the average delay to reinforcement from the onset of the initial links in DRT (T). According to DRT, conditioned reinforcing effectiveness is a function of the difference between T and the terminal-link delay to reinforcement. Preference is determined by relative conditioned reinforcing effectiveness, which increases as absolute terminal-link duration increases with the ratio held constant. T, like Log C, depends on both initial- and terminal-link durations, and serves as a comparator in determining preference. However, the models differ in the details of the comparator process, and the use of linear or logarithmically-scaled delays. DRT is also unable to predict a downturn in preference as absolute terminal-link duration increases (see Grace, 2004).

We calculated the criterion as the average of the log initial and terminal-link delays in each session, but a more realistic assumption would be to presume that there is a specific mechanism for updating the criterion. Perhaps the simplest way to accomplish this is to calculate the criterion as an exponentially weighted moving average (Killeen, 1981) of the delays between reward-correlated stimulus transitions:

$$\log C_{N+1} = \beta(\log D_N) + (1 - \beta)\log C_N \quad (6)$$

where log C_N and log $C_{N\text{-}1}$ are the criterion values after stimulus transitions N and N-1, respectively, log D_N is the Nth stimulus-transition delay, and β is a parameter that determines how much weight to give to the most recent delay. Note that N does not correspond to cycle number, as in Equation 3, because the criterion is updated twice per cycle—first after terminal-link entry (i.e., the initial-link -> terminal-link onset delay), and then again after food delivery (i.e., the terminal link -> food delay). With the addition of Equation 6, the model can be applied to situations in which the criterion might shift within sessions, for example, in which the initial- or terminal-link schedules are changed during a session. A goal for future research will be to explore how preference adapts in such situations, and whether Equation 6 is adequate as a representation of the criterion in the model.

Finally, it is worth noting that the assumption in the ExtDM that a fraction of the change in response strength during a session carries over to the next session provides a natural explanation for spontaneous recovery in choice behavior. For example, Mazur (1995, 1996) found that when the proportion of reinforcers delivered by an alternative was changed midway through a session (e.g., from 10% to 90%), pigeons' response allocation would shift (e.g., from 10% to 75%), but at the start of the next session would have reverted to an earlier percentage (e.g., 45%). Mazur proposed that this effect, which resembles spontaneous recovery, could be accounted for by assuming that the response strengths at the start of a session were determined by a weighted average of the several previous sessions. The ExtDM can predict the same result through a different but arguably simpler mechanism.

Thus, our results show that the extended version of Grace and McLean's (2006) decision model can be applied effectively to a situation in which terminal-link delays change systematically across sessions. The model's ability to predict the terminal-link effect in the present data, combined with Christensen and Grace's (2008) demonstration that the model can account for the initial-link effect, raises the possibility that the decision model may eventually provide a unified account of choice in concurrent chains under both acquisition and steady-state conditions.

REFERENCES

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior, 22,* 231–242.

Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior, 32*, 269–281.

Baum, W. M., & Rachlin, H. (1969). Choice as time allocation. *Journal of the Experimental Analysis of Behavior, 12*, 861–874.

Christensen, D. R., & Grace, R. C. (2008). Rapid acquisition in concurrent chains: effects of initial-link duration. *Behavioural Processes, 78*, 217–223.

Davison, M. (1983). Bias and sensitivity to reinforcement in a concurrent chain schedule. *Journal of the Experimental Analysis of Behavior, 40*, 15–34.

Davison, M., & Baum, W. M. (2000). Choice in a variable environment: Every reinforcer counts. *Journal of the Experimental Analysis of Behavior, 74*, 1–24.

Davison, M., & McCarthy, D. (1988). *The Matching Law: A research review.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Fantino, E. (1969). Choice and rate of reinforcement. *Journal of the Experimental Analysis of Behavior, 12*, 723–730.

Fantino, E., Preston, R., & Dunn, R. (1993). Delay reduction: Current status. *Journal of the Experimental Analysis of Behavior, 60*, 159–169.

Fantino, E., & Romanowich, P. (2007). The effect of conditioned reinforcement rate on choice: A review. *Journal of the Experimental Analysis of Behavior, 87*, 409–421.

Field, D., Tonneau, F., Ahearn, W., & Hineline, P. (1996). Preference between variable-ratio and fixed-ratio schedules: Local and extended relations. *Journal of the Experimental Analysis of Behavior, 66*, 283–295.

Fleshler, M., & Hoffman, M. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior, 5*, 529–530.

Gentry, D. G., & Marr, J. M. (1980). Choice and reinforcement delay. *Journal of the Experimental Analysis of Behavior, 33*, 27–37.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84*, 279–325.

Grace, R. C. (1994). A contextual model of concurrent chains choice. *Journal of the Experimental Analysis of Behavior, 61*, 113–129.

Grace, R. C. (1996). Choice between fixed and variable delays to reinforcement in the adjusting-delay procedure and concurrent chains. *Journal of Experimental Psychology: Animal Behavior Processes, 22*, 362–383.

Grace, R. C. (2002). Acquisition of preference in concurrent chains: Comparing linear-operator and memory-representational models. *Journal of Experimental Psychology: Animal Behavior Processes, 28*, 257–276.

Grace, R. C. (2004). Temporal context in concurrent chains: I. Terminal-link duration. *Journal of the Experimental Analysis of Behavior, 81*, 215–237.

Grace, R. C., & Bragason, O. (2004). Does the terminal-link effect depend on duration or reinforcement rate? *Behavioural Processes, 67*, 67–79.

Grace, R. C., Bragason, O., & McLean, A. P. (2003). Rapid acquisition of preference in concurrent chains. *Journal of the Experimental Analysis of Behavior, 80*, 235–252.

Grace, R. C., & McLean, A. (2006). Rapid acquisition in concurrent chains: Evidence for a decision model. *Journal of the Experimental Analysis of Behavior, 85*, 181–202.

Herrnstein, R. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior, 4*, 267–272.

Hunter, I., & Davison, M. C. (1985). Determination of a behavioral transfer function: White-noise analysis of session-to-session response-ratio dynamics on current VI VI schedules. *Journal of the Experimental Analysis of Behavior, 43*, 43–59.

Killeen, P. R. (1981). Incentive theory. *Nebraska Symposium on Motivation, 29.* Lincoln, University of Nebraska Press.

Kyonka, E. G. E., & Grace, R. C. (2007). Rapid acquisition of choice and timing in pigeons. *Journal of Experimental Psychology. Animal Behavior Processes, 33*, 392–408.

Kyonka, E. G. E., & Grace, R. C. (2008). Rapid acquisition of preference in concurrent chains when alternatives differ on multiple dimensions of reinforcement. *Journal of the Experimental Analysis of Behavior, 89*, 49–69.

MacEwen, D. (1972). The effects of terminal-link fixed interval and variable-interval schedules on responding under concurrent chained schedules. *Journal of the Experimental Analysis of Behavior, 18*, 253–261.

Mazur, J. E. (1984). Tests for an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology: Animal Behavior Processes, 10*, 426–351.

Mazur, J. E. (1992). Choice behaviour in transition: Development of preference with ratio and interval schedules. *Journal of Experimental Psychology: Animal Behavior Processes, 18*, 364–378.

Mazur, J. E. (1995). Development of preference and spontaneous recovery in choice behavior with concurrent variable-interval schedules. *Animal Learning and Behavior, 23*, 93–103.

Mazur, J. E. (1996). Past experience, recency, and spontaneous recovery in choice behavior. *Animal Learning and Behavior, 24*, 1–10.

Mazur, J. E. (2001). Hyperbolic value addition and general models of animal choice. *Psychological Review, 108*, 96–112.

Mazur, J. E., Blake, N., & McManus, C. (2001). Transitional choice behavior in concurrent-chain schedules. *Behavioural Processes, 53*, 171–180.

Navarick, D. J., & Fantino, E. (1976). Self-control and general models of choice. *Journal of Experimental Psychology: Animal Behavior Processes, 2*, 75–87.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109*, 472–491.

Schofield, G., & Davison, M. (1997). Non-stable concurrent choice in pigeons. *Journal of the Experimental Analysis of Behavior, 68*, 219–232.

Shull, R. L., Spear, D. J., & Bryson, A. E. (1981). Delay or rate of food delivery as a determiner of response rate. *Journal of the Experimental Analysis of Behavior, 35*, 129–143.

Spetch, M. L., & Dunn, R. (1987). Choice between reliable and unreliable outcomes Mixed percentages – reinforcement in concurrent chains *Journal of the Experimental Analysis of Behavior, 47*, 57–72.

Williams, B., & Fantino, E. (1978). Effects on choice of reinforcement delay and conditioned reinforcement. *Journal of the Experimental Analysis of Behavior, 29*, 77–86.